About Big Red 3 at Indiana University

On this page:

- System overview
- System access
- HPC software
- Set up your user environment
- File storage options
- Work with data containing PHI
- Run jobs on Big Red 3
- Acknowledge grant support
- Get help

System overview

Big Red 3 is a Cray XC40 supercomputer dedicated to researchers, scholars, and artists with large-scale, compute-intensive applications that can take advantage of the system's extreme processing capability and high-bandwidth network topology. Big Red 3 supports programs at the highest level of the university, including the Grand Challenges program.

Featuring 930 dual-socket compute nodes equipped with Intel Haswell Xeon processors (22,464 compute cores), Big Red 3 has a theoretical peak performance (Rpeak) of 934 trillion floating-point operations per second (934 teraFLOPS).

Big Red 3 runs a proprietary variant of Linux called Cray Linux Environment (CLE). In CLE, compute elements run a lightweight kernel called Compute Node Linux (CNL), and the service nodes run SUSE Enterprise Linux Server (SLES). Big Red 3 uses Slurm to coordinate resource management and job scheduling.

System access

IU students, faculty, staff, and affiliates can create Big Red 3 accounts using the instructions in Get additional IU computing accounts (achr). Grand Challenges users who create Big Red 3 accounts can use IU HPC Projects ⬀ (https://hpcprojects.uits.iu.edu/) to request access to the Grand Challenge partition; for instructions, see Request and manage access to specialized high performance computing (HPC) resources (bgmo).

Once your account is created, you can use any SSH2 (aelc) client to access `bigred3.uits.iu.edu`. Log in with your IU username and passphrase (beml).

> **Notes:**
>
> - Two-factor authentication using Two-Step Login (Duo) (beum) is required for access to the login nodes on IU research supercomputers (alde), and for SCP and SFTP file transfers to those systems. SSH public key authentication remains an option for researchers who submit the "SSH public key authentication to HPS systems" agreement ⬀ (https://hpceverywhere.iu.edu/agree#/agree) (log into HPC *everywhere* (aquv) using your IU username and passphrase (beml)), in which you agree to set a passphrase on your private key when you generate your key pair. If you have questions about how two-factor authentication may impact your workflows, contact the UITS Research Applications and Deep Learning team ⬀ (https://mailform.kb.iu.edu/email.php?cid=1449). For help, see Get started with Two-Step Login (Duo) at IU (bfgm) and Help for Two-Step Login (Duo) (aluu).
>
> - For enhanced security, SSH connections that have been idle for 60 minutes will be disconnected. To protect your data from misuse, remember to log off or lock your computer whenever you leave it.

- The scheduled monthly maintenance window for IU's high performance computing systems is the second Sunday of each month, 7am-7pm.

HPC software

The Research Applications and Deep Learning (RADL) group, within the Research Technologies ⎋ (https://rt.iu.edu) division of UITS (ahaw), maintains and supports the high performance computing (HPC) software on IU's research supercomputers (alde). To see which applications are available on a particular system, log into the system, and then, on the command line, enter `module avail` .

For information on requesting software, see Software requests (avkh#software) in Policies regarding UITS research systems (avkh).

Set up your user environment

The IU research supercomputers use module-based environment management systems that provide a convenient method for dynamically customizing your software environment. Big Red 3 uses the Modules module management system.

For more, see Use modules to manage your software environment on IU research supercomputers (bcwy).

Big Red 3 provides programming environments for the Cray, Intel, PGI, and GNU Compiler Collections (GCC) compilers. For information about using these compiler suites, see Compile programs on Big Red 3 at IU (bdnv).

File storage options

For file storage information, see Available access to allocated and short-term storage capacity on IU's research systems (avkm).

Before storing data on this system, make sure you understand the information in the Work with data containing PHI section (below).

Work with data containing PHI

The Health Insurance Portability and Accountability Act of 1996 (HIPAA (ayyy)) established rules protecting the privacy and security of individually identifiable health information. The HIPAA Privacy Rule ⎋ (https://www.hhs.gov/hipaa/for-professionals/privacy/index.html) and Security Rule ⎋ (https://www.hhs.gov/hipaa/for-professionals/security/index.html) set national standards requiring organizations and individuals to implement certain administrative ⎋ (https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/securityrule/adminsafeguards.pdf), physical ⎋ (https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/securityrule/physsafeguards.pdf), and technical ⎋ (https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/securityrule/techsafeguards.pdf) safeguards to maintain the confidentiality, integrity, and availability of protected health information (PHI (ayyz)).

This UITS system or service meets certain requirements established in the HIPAA Security Rule ⎋ (https://www.hhs.gov/hipaa/for-professionals/security/index.html) thereby enabling its use for work involving data that contain protected health information (PHI (ayyz)). However, using this system or service does not fulfill your legal responsibilities for protecting the privacy and security of data that contain PHI. You may use this system or service for work involving data that contain PHI only if you institute additional administrative ⎋ (https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/securityrule/adminsafeguards.pdf),

physical 🗗
(https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/securityrule/physsafeguards.pdf), and
technical 🗗
(https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/administrative/securityrule/techsafeguards.pdf)
safeguards that complement those UITS already has in place.

For more, see Your legal responsibilities for protecting data containing protected health information (PHI) when using
UITS Research Technologies systems and services (ayzm).

> **Note:**
>
> Although PHI (ayyz) is classified as Critical data 🗗 (https://datamanagement.iu.edu/data-classifications/critical-
> data/index.html), other types of institutional data classified as Critical are not permitted on Research
> Technologies 🗗 (https://rt.iu.edu) systems. For help determining which institutional data elements classified as
> Critical are considered PHI, see About protected health information (PHI) data elements in the classifications of
> institutional data (bdtx).

If you have questions about securing HIPAA-regulated research data at IU, email `securemyresearch@iu.edu` .
SecureMyResearch provides self-service resources and one-on-one consulting to help IU researchers, faculty, and staff
meet cybersecurity and compliance requirements for processing, storing, and sharing regulated and unregulated
research data; for more, see About SecureMyResearch (bgin). To learn more about properly ensuring the safe handling
of PHI (ayyz) on UITS systems, see the UITS IT Training video Securing HIPAA Workflows on UITS Systems 🗗
(https://iu.mediaspace.kaltura.com/media/Securing+HIPAA+Workflows+on+UITS+Systems/1_8vw13atl). To learn
about division of responsibilities for securing PHI, see Shared responsibility model for securing PHI on UITS systems
(aqqk).

Run jobs on Big Red 3

Big Red 3 uses the Slurm workload manager for resource management and job scheduling; see Use Slurm to submit
and manage jobs on high performance computing systems (awrz).

In Slurm, compute resources are grouped into logical sets called partitions, which are essentially job queues. To view
details about available partitions and nodes, use the `sinfo` command; for more about using `sinfo` , see the View
partition and node information (awrz#sinfo) section of Use Slurm to submit and manage jobs on high performance
computing systems (awrz).

> **Notes:**
>
> - Big Red 3 jobs are allocated memory based on task count (`--ntasks-per-node`). Each task receives 1.2 GB of
>   memory (for example, four tasks will receive 4.8 GB, and so on). If needed, users may request additional
>   memory using the `--mem` flag. The maximum memory you can request on a node is 58 GB. For example, in a
>   batch script:
>
>   ```
>   #SBATCH –mem=58G
>   ```
>
>   For an interactive job request:
>
>   ```
>   srun -p general -N 1 --ntasks-per-node=4 --mem=4G --time=1:00:00 --pty bash
>   ```
>
> - Each Big Red 3 compute node has 48 hyper-threaded CPU cores and 24 physical CPU cores. Users can
>   compute nodes among their own jobs when not requesting all the cores on a node. If none of the jobs has
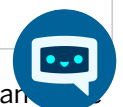
unique resource constraints, such as special memory requirements, Slurm automatically stacks up to 24 jobs on one compute node. For example, if you submit four jobs, each requesting 12 cores, Slurm will place them on the same node. To prevent this behavior, you can request all the cores on a node by adding the following directives to your Slurm job script:

```
#SBATCH --nodes=1
#SBATCH --ntasks-per-node=48
```

If you request all 48 cores on a node but want to use fewer cores for your run, use the `-n` flag to specify that in your `srun` command; for example:

```
srun -n 24 a.out
```

Otherwise, `srun` will launch the application with 48 cores.

Acknowledge grant support

The Indiana University cyberinfrastructure, managed by the Research Technologies ☐ (https://rt.iu.edu) division of UITS (ahaw), is supported by funding from several grants, each of which requires you to acknowledge its support in all presentations and published works stemming from research it has helped to fund. Conscientious acknowledgment of support from past grants also enhances the chances of IU's research community securing funding from grants in the future. For the acknowledgment statement(s) required for scholarly printed works, web pages, talks, online publications, and other presentations that make use of this and/or other grant-funded systems at IU, see Sources of funding to acknowledge in published work if you use IU's research cyberinfrastructure (anwt)

Get help

Support for IU research supercomputers, software, and services is provided by various teams within the Research Technologies ☐ (https://rt.iu.edu) division of UITS (ahaw).

- If you have a system-specific question, contact the High Performance Systems (HPS) team ☐ (https://mailform.kb.iu.edu/email.php?cid=26).

- If you have a programming question about compilers, scientific/numerical libraries, or debuggers, contact the UITS Research Applications and Deep Learning team ☐ (https://mailform.kb.iu.edu/email.php?cid=1449).

For general questions about research computing at IU, contact UITS Research Technologies ☐ (https://rt.iu.edu/support/index.html).

For more options, see Research computing support at IU (alfa).

This is document aoku *in the Knowledge Base.*
Last modified on 2021-11-11 13:25:32.

v.facebook.com/iu.uits) (https://twitter.com/insideuits) (https://www.instagram.com/uitsatiu/) (http://www.youtube.cc

Support & More

General support or comments:

- Email the UITS Support Center (https://mailform.kb.iu.edu/email.php?cid=65)

Research technology support:

- Contact Research Technologies (https://rt.iu.edu/support/index.html)

Chat with a consultant (http://ithelplive.iu.edu)

AskIU (https://mailform.kb.iu.edu/email.php?cid=1061)

One.IU (https://one.iu.edu)

Navigation

Home (/)

Menus (/d/menu)

About us (/d/line)

IT@IU

About IT (http://it.iu.edu/about/)

IT staff (https://uits.iu.edu/staff)

Jobs in IT (http://it.iu.edu/jobs/)

OVPIT (http://it.iu.edu/ovpit/)

Ψ

Accessibility (https://accessibility.iu.edu/assistance/)

Privacy (/d/priv)

Copyright (https://www.iu.edu/copyright/index.html) © 2022 The Trustees of Indiana University (https://iu.edu)